

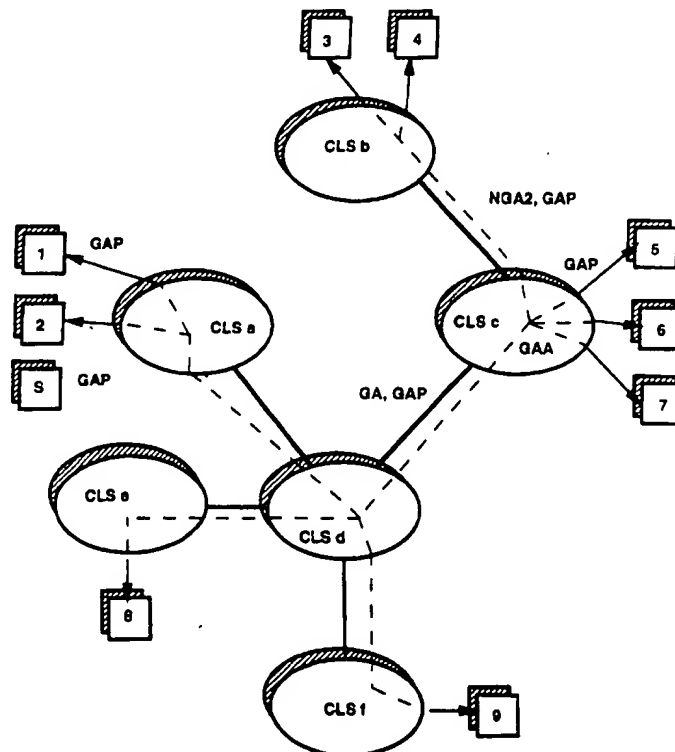


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶: H04L 12/56	A1	(11) International Publication Number: WO 96/38961 (43) International Publication Date: 5 December 1996 (05.12.96)
(21) International Application Number: PCT/SE96/00682 (22) International Filing Date: 28 May 1996 (28.05.96) (30) Priority Data: 9501999-8 31 May 1995 (31.05.95) SE (71) Applicant (for all designated States except US): TELIA AB [SE/SE]; S-123 86 Farsta (SE). (72) Inventor; and (75) Inventor/Applicant (for US only): KAVAK, Nail [SE/SE]; Myrstuguvägen 359, S-143 32 Vårby (SE). (74) Agent: KARLSSON, Berne; Telia Research AB, Rudsjöterrassen 2, S-136 80 Haninge (SE).		(81) Designated States: NO, US, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: METHOD AND TRANSMISSION SYSTEM RELATED TO MULTICASTING**(57) Abstract**

The present invention relates to multicasting in an ATM based public data network. Multicasting can be defined as the ability to send a single message to multiple recipients at different locations. A method of multicasting, of the present invention, operates in a packet switched data transmission system typically an ATM broadband data service system. A data packet, intended for transmission to a plurality of addresses, is copied at each node of the system to which one, or more, subscribers, to which said data packet is addressed, are connected. A copy of said data packet is transmitted to each addressed subscriber and to adjacent nodes to which addressed subscribers are connected. This procedure is continued until all addressed subscribers have received a copy of the data packet.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	L	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

Method and transmission system related to
multicasting

The present invention relates to a method of multicasting in a
packet switched telecommunications system, mor particularly an ATM
5 based public data network, and a telecommunications network adapted
to facilitate multicasting.

It frequently happens that a subscriber to a data network needs to
send the same data to several different addresses, preferably at
10 substantially the same time. This facility is known as multicasting,
and various techniques are known for its implementation. Examples of
telecommunications applications that support multicasting are:

- group communication
- 15 - virtual private networks
- video and audio conferencing
- discovering routes in remote bridging applications
- routing information exchanges
- file server recovery
- 20 - address resolution to discover physical addresses
- distributed simulations and games, such as tank
battle simulations involving several geographically
separated participants.

25 Multicasting can be defined as the ability to send a single
message to multiple recipients at different locations.

Most existing data communication networks, e.g. Internet, SMDS, ATM
LAN emulation, support multicasting/broadcasting. Broadband data
30 services, (BDS) must also support multicasting if the service is to
be attractive to users. Group addresses are used to support
multicasting applications in these networks.

Typically a BDS is implemented at various ATM switches. The
35 connectionless server functions, (CLSs), may be installed within, or
associated with, ATM switches. CLSs may not be connected in all
switching nodes. The CLSs may constitute mesh, ring, star, or tree
networks. The CLSs control operation of the individual networks and
switch packets within networks and between networks.

Typically, in ATM, an information packet, or cell, comprises five bytes of header data and 48 bytes of user data. The header contains data that identifies the cell, a logical address that identifies routing, forward error correction bits, plus bits for priority handling and network management functions. Forward error correction is only applied to the header. All cells of a virtual connection follow the same path through the network, which is determined during call set-up. There are no fixed time slots so that any user can access the transmission medium whenever an empty cell is available.

10 ATM is capable of operating at bit rates of 155.52 and 622.08 Mbit/s.

Where a BDS, operating on a data packet transmission technique, such as ATM, includes a multiplicity of switching nodes, the transmission of a plurality of identical data packets from node A, via nodes B, C, D and E, to subscribers served from node E, wastes considerable traffic capacity. The same information is transmitted between the same nodes many times. The present invention is directed to reducing this redundant transmission.

20

The present invention uses "agents", in its implementation, which are object oriented programmed entities including both data and operating code, and having the capacity to "negotiate" with other agents, or entities, e.g. subscribers.

OF THE INVENTION

In common with general practice in the telecommunications field, this specification makes use of abbreviations. To assist the reader, the specific description starts with a glossary of abbreviations used in the specification.

30

The present invention can, of course, be used with advantage in any packet switched telecommunications network employing a mesh, ring, or tree transmission architecture in which transmission paths are interconnected by switching nodes, some of which are CLS nodes, i.e. switching nodes having a CLS function. However, the invention has actually been implemented using an experimental BDS system developed by the applicant. For the sake of completeness a brief description of the experimental BDS system, is included in the present specification.

35

SUBSTITUTE SHEET

According to a first aspect of the present invention, there is provided a method of multicasting in a packet switched data transmission system, having a plurality of switching nodes, each serving a plurality of subscribers, characterised in that a data packet, intended for transmission to a plurality of addresses, is copied at each node to which one, or more, subscribers, to which said data packet is addressed, are connected, a copy of said data packet being transmitted to each addressed subscriber and to the next node to which addressed subscribers are connected, until all addressed subscribers have received a copy of said data packet.

Preferably a CLS node transmits one and only one copy of said data packet to each CLS node:

- 15 - to which it is directly connected; and
- to which addressed subscribers are connected.

Alternatively a CLS node transmits one and only one copy of said data packet to each CLS node:

- to which it is directly connected, and
- to which either:
 - 25 - addressed subscribers; or
 - a CLS node connected to addressed subscribers;
- 30 are connected.

Each multicast data packet may include a group address.

A multicast data packet transmitted from a first CLS node and destined for a plurality of subscribers is transmitted from said first CLS node in a data packet which encapsulates the group address.

Preferably data packets are transferred along paths in the form of a tree, packets only being copied where branches of the tree diverge, and in that the structure of said tree is determined so that system usage is minimised, subject to the constraint that transmission
5 delays to each group address are kept within a preset limit.

Said paths may constitute a Steiner tree.

According to a second aspect of the present invention, there is
10 provided a packet switched data transmission system, having a plurality of switching nodes, each serving a plurality of subscribers, characterised in that each node includes copying means for duplicating data packets, intended for transmission to a plurality of addresses, and transmission means adapted to transmit
15 copies of multi-addressed data packets to each addressed subscriber and to the next node to which addressed subscribers are connected.

Preferably said transmission means are adapted to transmit one and only one copy of said data packet to each CLS node;

20

- to which its CLS node is directly connected; and

- to which addressed subscribers are connected.

25 Alternatively said transmission means are adapted to transmit one and only one copy of said data packet to each CLS node:

- to which its CLS node is directly connected; and

30

- to which either:

- addressed subscribers; or

- a CLS node connected to addressed subscribers;

35

- are connected.

Each multicast data packet may include a group address.

Said packet switched data transmission system may include at least one ATM based public data network.

5 Preferably at least one CLS node, in said packet switched data transmission system, includes group addressing means, having, for each group address, a data base of all individual addresses comprising a group address, and in that said group addressing means
10 is adapted to control a multicast to a group address and to completely resolve a group address.

Alternatively one and only one CLS node, in said packet switched data transmission system, includes group addressing means, having,
15 for each group address, a data base of all individual addresses comprising a group address, said group addressing means adapted to control a multicast to a group address and capable of completely resolving a group address.

20 Alternatively one and only one CLS node, in said packet switched data transmission system, includes a first group addressing means, said first group addressing means being adapted to control a multicast to a group address, and to partially resolve a group address, in that each of a plurality of CLS nodes has a second group addressing means
25 adapted to complete resolution of a group address which has been partially resolved by said first group addressing means, in that only one first group addressing means is permitted per group address and in that more than one second group addressing means is permitted per group address.

30 Embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings in which:

Figure 1 illustrates the composition of an ATM address.

35 Figure 2 illustrates a centralised group address agent structure.

Figure 3 illustrates a path tree based on a source/destination routing calculation.

Figure 4 illustrates the protocol architecture for a
5 connectionless service using CL servers.

Figure 5 illustrates cell processing in a CL server.

- a) - enter/activate the E.164 address
 - 10 - check the qos
 - establish a new route
 - transfer the cell

 - b) - check the incoming VCI
 - 15 - consult the routing table
 - retrieve the outgoing VCI
 - transfer the cell

 - c) - check the incoming VPI/VCI
 - 20 - consult the routing table
 - retrieve the outgoing VPI/VCI
 - transfer the cell
 - delete (inactivate) inc/out VPI/VCI
- 25 A glossary of the abbreviations used in this specification is set out below:

- AAL3 - ATM adaptation layer, type 3
- AAL4 - ATM adaptation layer, type 4
- 30 AAL5 - ATM adaptation layer, type 5
- AC - access code
- ATM - asynchronous transmission mode
- B-ISDN - broadband integrated services digital network
- BDS - broadband data service
- 35 BOM - beginning of message
- CBDS - connectionless broadband data service
- CC - country code
- CL - connectionless
- CLNAP - connectionless network access protocol

- CLNAP-PDU - group addressed data packet, more generally a CLNAP protocol data unit.
- CLNIP - connectionless network interface protocol
- CLS - connectionless server
- 5 CLSF - connectionless service function
- cnf - extension for confirm messages
- COM - continuation of message
- CPE - customer premises equipment
- DG - designated gateway
- 10 EOM - end of message
- FTP - file transfer protocol
- GA - group address
- GAA- group address agent
- GAP- group addressed data packet
- 15 IP- internet protocol
- LAN- local area network
- LEC - local exchange code
- ME - mapping entity
- NDC - national destination code
- 20 NGAA - nested group address agent
- NNI - network node interface
- OSPF - open shortest path first
- PDU - packet data unit
- PVC - permanent virtual circuit
- 25 QoS - quality of service
- RD - Routing domain
- req - extension for request message
- SMDS - switched multi-megabit data service
- SN - subscriber number
- 30 SSM - single segment message
- SVC - signalling virtual circuit
- UNI - user network interface
- VCI - virtual circuit identifier
- VP - virtual path
- 35 VPI- virtual path identifier

A multicast information transfer allows a subscriber to send data packets to a packet switched data system which are transferred, by the system, to several recipients. The present invention employs group

addressing to achieve multicasting, and is described in relation to an ATM based public data service. It will be obvious, to those skilled in the art, how to adapt the present invention to other packet switched data transfer systems.

5

The system is required to deliver one and only one copy of a group addressed data packet (GAP) across each of the connectionless access interfaces associated with the individual addressees represented by the group address. Any recipient of a GAP may make use of the
10 destination group address, carried by that GAP, to multicast to addressees of the GAP, (except himself). Non-members of a group identified by a group address (GA) may send GAPs to that group. The service provider is responsible for assigning group addresses and ensuring that each GA identifies uniquely only one set of individual
15 addresses. Within the network, the group addressed data packets are mapped into AAL5 streams.

The address structure used in the BDS, of the present embodiment of the invention, takes the form:

20

<C|E><E.164 number>

The address type, C, or E, is a hexadecimal digit which identifies the address as an individual address, or a group address. For example, C 4687075565_H represent a BDS individual address, and E4687075565_H represents a BDS group address. The composition of ATM
25 addresses is illustrated in Figure 1. The ATM address is first split into address type and E-164 number. The E.164 number can then be broken down into:

- country code, CC, (1 to 3 digits);
- 30 - national destination code, NDC;
- subscribers number, SN (8 digits);
- routing domain, RD, (1 to 2 digits);
- area (3 to 4 digits);
- local exchange code, LEC, (3 digits); and
- 35 - access code, AC, (5 digits);

as shown in Figure 1.

A group address agent (GSA) serves as the centralised administrator for a group address. The GAA interacts with a subscriber to handle requests for the creation, modification, or deletion, of group addresses and, provides information on membership
5 of a group address. Group addresses can be stored within the BDS in three ways:

- centralised storage;
- replicated storage; and
- 10 - nested centralised storage.

In the centralised approach to storage of group addresses one of the connectionless service nodes administers the assignment, deletion and modification of group addresses, and inclusion, addition and
15 deletion of individual addresses to the group, based on the instructions of the user. From the numbering point of view, the GAA assigns, to a group, a globally unique group address which pertains to the GAA's domain. The GAA performs a complete address resolution function. This approach has advantages and disadvantages.

20

The advantages of the centralised approach are;

- it creates an overall view of an address group within a BDS system;
- 25 - easier administration and maintenance;
- no need to introduce special measure to ensure consistency; and
- it is easy to migrate to a nested group addressing methodology.

The disadvantages of the centralised approach are:

30

- transit delay is increased;
- the load on a single address agent is increased;
- the network load is increased; and
- a lack of robustness to single point failure, (which can be
35 ameliorated by redundant agent implementation).

In the replicated, i.e. distributed, data base approach, the resolution of group addresses requires each network to have its own data base which must hold all the information relating to all group

addresses defined on all networks, regardless of their location. This approach provides minimum transit delay since each copy is directly routed towards its destination after resolution of the group address within the originating network. It is robust because, if the network fails, remaining networks may still resolve the group addressed packets. Furthermore, fewer packets are generated. On the other hand, this methodology increases network loading since the total multicasting occurs at the originating network. Maintaining consistency of up to date information is also difficult. For these reasons this approach does not form part of the present invention.

Nested group addressing is used in conjunction with a centralised database scheme. One group address agent manages each individual group address, as previously described with reference to the centralised approach. Nested group addressing can be used to avoid sending repetitive information to individual members of a group who are accessed via the same network and employs an NGAA. While only one GAA is allowed per group address, one NGAA may be dedicated to serve each separate network making up the BDS. There may be more than one NGAA per group address. Each network having one, or more, members of the GA, may have NGAA functionality. Thus, in the nested group address approach, the GAA performs a partial address resolution function. Nested group addressing is an efficient mechanism for transport of group addressed packets. It results in lower network loading because only one copy of the data needs to be sent to each of the networks involved. However, it also causes increased transit delay because the packets need to be routed to the central multicasting network. Another problem with this approach is lack of robustness, single point failure can cause failure to deliver, or administer, group address packets.

The centralised database approach is favoured when the number of customers is reasonably low. However, when the number of customers increases, the centralised approach is upgraded to a nested group addressing scheme.

It is assumed that the ATM address of the GAA will be well known by all subscribers to the BDS. A user can thus set up a control channel to the GAA to send requests on registration/deregistration of individual addresses and receive replies on corresponding group

SUBSTITUTE SHEET

addresses. The channel used for issuing multicast packets (data channel) to the GAA may, or may not, be the same as the control channel. There may, however, be some advantage to separating the control channel from the data channel, as the QoS and bandwidth requirements may be different for the different channels. Preferably, the control channel can be released after a pre-defined inactivity timeout has elapsed.

10 The protocol for registration of a group address is as outlined below.

All individual addresses that belong to the same group are registered to a central GAA entity. A designated subscriber is an end user who is authorised to create a new group address. A designated subscriber is not entitled to create a new group address which excludes himself.

15 The following messages can be exchanged between a designated subscriber and the GAA entity:

```
Registersubs_req (IndATMaddr1; IndATMaddr2:
.....IndATMaddr3, transaction_ID)

20 Registersubs_cnf (GroupATMaddress, transaction_ID,
failure/success)
```

Non-group members can send packets to a group. The confirm messages are only issued to the designated subscriber. (Optionally, the confirm message can be sent to all members of the group for security reasons). The transaction_ID in the confirm message should be the same as in the registration message. This ID is used for local purposes, (e.g. validity of the confirms). A given group address can have only one GAA.

30

The designated subscriber can, through interaction with the GAA entity, delete, or add, a new group member to an existing group, of which he is a member. The designated subscriber can even query the GAA to resolve a group address. This function may be usefully employed to check the current state of the group membership.

35

To prevent, among other things, non-group members becoming group members, without notification to other group members, only a designated subscriber is permitted to initiate the following operational commands:

```
5      Deletesubs_req (GroupATMaddress, IndATMaddr1:
      IndATMaddr2 : .... IndATMaddr3, transaction_ID)

      Deletesubs_cnf (GroupATMaddress, transaction_ID,
10     failure/success)

      Addsubs_req (GroupATMaddress, IndATMaddr1:
      IndATMaddr2:..... IndATMaddr3, transaction_ID),

15     Addsubs_cnf (GroupATMaddress, transaction_ID,
      failure/success)

      Resolvegroup_req (GroupATMaddress, transaction_ID)

20     Resolvegroup_cnf (IndATMaddr1: IndATMaddr2:
      ... IndATMaddr3, transaction_ID)
```

25 A group address with no members is automatically deleted by the
GAA.

Multicasting packets are issued in accordance with the following protocol:

- 30 - source S sends a packet with destination address GA to the
nearest connectionless node N, say;
- N recognises GA as a group address and encapsulates the original
packet in a new envelope with source address (S,GA);
- N forwards the packet to the central GAA entity;
- the GAA receives the packet and resolves the GA using the
- 35 centrally managed address table;
- the GAA issues as many packets as are necessary for transmission
to group members within the group; and
- in order to transfer the upicast packets, the GAA entity
computes the shortest paths to all members of the group.

SUBSTITUTE SHEET

It should be noted that encapsulation means that the user data unit, when addressed to a remote node, is encapsulated as the payload of a new data unit whose header contains proper addressing information needed for it to reach its destination.

The superposition of all the shortest paths computed by the GAA entity constitute a multicast tree.

Detailed operation of this protocol can be better understood by considering an example illustrated by Figure 2. The DG identifies members designated 1 to 9, in Figure 2. CLSa contains the GAA for the GA. Operation of the centralised scheme protocol proceeds as follows:

- user S, connected to CLSa, originates a CLNAP-PDU with DA = GA(GAP);
- this PDU is routed to CLSc, which contains the GAA for the GA;
- CLSc resolves the GA into individual addresses of all members, (i.e. 1 through 9);
- CLSc delivers the GAP directly to its own members, (5,6,7) and sends to each of the external members a CLNIP-PDU carrying the individual destination address of the member;
- no packet is sent back to the sender.

Multicast, i.e. group addressed packets, are transferred encapsulated within the data field of a CLNAP-PDU. The GAA can transfer the individual packets either by point-to-point connections, or point-multipoint connection, where the GAA is the root of a tree and individual members are leaves of the tree.

The multicast packets are replicated from the GAA to all other recipients, including network nodes, i.e. ATM nodes as well as connectionless servers. The multicast packets are transferred via paths which form a tree structure. The tree is rooted at the GAA and the branches terminate at subscribers who are members of the group. The packets are replicated when, and only when, two branches of the path diverge. Postponing replication of the packets, in this way, conserves network bandwidth. If a single link leads to multiple group

members, only a single copy of the packet traverses the link, if necessary it will be replicated later. Since the multicast path is calculated as a tree at the GAA, the path taken by a multicast packet depends both on the location of the GAA in relation to the source and its multicast destination. This mechanism is called source/destination routing and should be contrasted with Open Shortest Path First (OSPF) routing, in which the route is based solely on destination. Taking account of the source when making routing decisions requires a lot more calculation but leads to very good paths in terms of network usage and delay to individual group members. In fact, the path, taken between packet source and any particular destination group member, is the least cost path available. By using source destination routing, as opposed to OSPF, a path similar to a minimal spanning tree is calculated. This type of path is referred to as a Steiner tree.

A typical path, corresponding to the centralised group address agent structure illustrated in Figure 2, is shown in Figure 3. The GAA, located in CSLC, is the root of the path. The branches of the path terminate at group members 1 to 9. It should be noted that the packets are replicated when, and only when, two branches of the path diverge. This results in a substantial reduction in copying, which means that there will be a substantial reduction in processing and reduced bandwidth consumption.

Routing calculations may be modified to take account of QoS required by particular applications, or group members, e.g. delay, throughput, reliability, cost, operator etc

The present invention has actually been realised on a prototype BDS based on the direct connectionless broadband data service (CBDS) model of the ITU/ETSI (1364) standard. This prototype will now be briefly described, to assist an understanding of how the present invention can be implemented. A Connectionless Service Function (CLS F) is installed within B-ISDN. The CLS F terminates CL protocols and routes CL-packets to their destinations, according to routing information included in CL-PDUs. The ATM connections between the user and the CLS F, as well as the connection between CLS Fs, can be either

SUBSTITUTE SHEET

PVC or SVC. The CL-servers can be installed at various ATM switches, as well as ATM cross connect nodes. The connectionless user sends data to a CL-server whose address and identity it knows, e.g. the closest one, to which it either signals or has semipermanent
5 connection already established. The CL-server then forwards the data to the destination user possibly by a route including other CL-servers. The CL-servers may be interconnected by a virtual overlay network consisting of several VPs with pre-allocated bandwidth resources. The choice of PVC, or SVC depends on the user traffic
10 characteristics and QoS requirements. The use of CL-servers within the ATM network will lead to a reduction of the number of VPs needed (as compared to the full VP mesh) and thus to a concentration of connectionless traffic on fewer VPs. By statistically multiplexing several sources on the same VP, burstiness can be reduced, although
15 the losses due to buffer overflow may be increased. Furthermore, the number of connections needed by each end point, for setup, is reduced to only one.

Figure 4 illustrates the protocol architecture of the CL-servers at
20 the user interface (UNI) as well as between the network nodes (NNI).

The CLSF terminates the B-ISDN CLNAP which includes functions for the mapping of connectionless protocols onto the connection-oriented ATM service by means of AAL3/4 entities. The CLNAP includes functions such
25 as routing, addressing, multicasting, QoS selection, carrier selection, etc.. Routing is based on the E.164 address information contained in the CLNAP-PDU header.

The CLNIP supports the CL-service between CL-servers inside a network operator domain as well as between network operator domains. CLNIP
30 provides for the transport of both encapsulated and non-encapsulated data units. Interworking functions between the CLNAP and the CLNIP are provided by a ME in CL-servers.

35 The CLNAP layer includes, among other things, functions for the routing and addressing of variable length CBDS packets transferred between one source and one, or more, destinations, without the establishment of any ATM connection by the user. It also supports multiprotocol encapsulation and the QoS parameters: transit delay,

cost and residual error probability. To achieve higher traffic concentration at CLSs, VPs between CLs can be configured so that they do not have full mesh connectivity. Alternatively, the CL-servers can be interconnected arbitrarily, or by means of other topological schemes such as hierarchical trees, buses, or rings.

The service provided by the prototype BDS is rather similar to the IEEE MAC service:

- 10 - BDS_UNITDATA.request (Source E.164, Destination E.164, User Data, QoS)
- BDS UNITDATA.indication (Source E.164, Destination E.164, User Data, QoS)

15 This service is implemented on a prototype BDS system which consists of a Fore System's ATM switch and 4 Sun/SPARC workstations. Each workstation is provided with a Fore System's ATM computer interface and a fibre connection to the switch. A workstation may be configured either as a terminal, or server. Implementation of BDS related

20 routines can be performed on the Fore System's host interface version 2.1.1.1, in a way that maximises use of existing Fore System's code for segmentation, reassembly, data structures, etc..

The terminal code essentially consists of an output routine and an

25 input routine that respectively construct and parse BDS PDU headers. On the output side, the routine invokes the Fore System's AAL5 protocol. Since this protocol requires a VPI/VCI pair on which to send, a table lookup has to be performed that yields such a pair corresponding to a given E.164 destination address. On the input side,

30 the AAL5 entity delivers a buffer together with a BDS PDU header which is parsed by the input routine and then forwarded to an appropriate higher level entity. At present, two upper layer protocols have been taken into account:

- 35 - IP for transparent use of current Internet applications such as FTP, virtual terminal (TELNET) etc.; and
- a "Raw socket" interface, that gives end users direct access to the BDS service described above.

SUBSTITUTE SHEET

Since the server operates on a per cell basis, the server code has to be placed at the lowest level possible over the Fore Systems's computer interface. When a cell header has been retrieved from the ATM cell input FIFO queue, the server code checks whether the VPI is one that has been designated as part of the BDS network (VPI = "cls"), as illustrated in Figure 5. If so, a table lookup is performed to find the appropriate outgoing VPI/VCI pair. For reasons of efficiency, table mappings are maintained per session (using timeouts), rather than per packet, since it can be assumed that an application entity will, in general, generate more than one packet in its interaction with applications entities on other systems. If a cell is part of a COM this amounts to indexing a table. More complicated procedures are required if a cell is a BOM, or EOM, see Figure 5.

CLAIMS

1. A method of multicasting in a packet switched data transmission system, having a plurality of switching nodes, each serving a plurality of subscribers, c h a r a c t e r i s e d in that a data packet, intended for transmission to a plurality of addresses, is copied at each node to which one, or more, subscribers, to which said data packet is addressed, are connected, a copy of said data packet being transmitted to each addressed subscriber and to the next node to which addressed subscribers are connected, until all addressed subscribers have received a copy of said data packet.

2. A method as claimed in claim 1, c h a r a c t e r i s e d in that a CLS node transmits one and only one copy of said data packet to each CLS node:

- to which it is directly connected; and
- to which addressed subscribers are connected.

3. A method as claimed in claim 1, c h a r a c t e r i s e d in that a CLS node transmits one and only one copy of said data packet to each CLS node:

- to which it is directly connected; and
- to which either:
 - addressed subscribers; or
 - a CLS node connected to addressed subscribers;are connected.

4. A method as claimed in any previous claim,
c h a r a c t e r i s e d in that each multicast data packet
includes a group address.

5 5. A method as claimed in any previous claim,
c h a r a c t e r i s e d in that said packet switched data
transmission system includes at least one ATM based public data
network.

10 6. A method as claimed in any previous claim,
c h a r a c t e r i s e d in that at least one GLS node, in said
packet switched data transmission system, includes group addressing
means, having, for each group address, a data base of all
individual addresses comprising a group address, and in that said
15 group addressing means is adapted to control a multicast to a group
address and to completely resolve a group address.

7. A method as claimed in claim 5,
c h a r a c t e r i s e d in that one and only one CLS node, in said
20 packet switched data transmission system, includes group addressing
means, having, for each group address, a data base of all individual
addresses comprising a group address, said group addressing means
adapted to control a multicast to a group address and capable of
completely resolving a group address.

25 8. A method as claimed in claim 5,
c h a r a c t e r i s e d in that one and only one CLS node, in said
packet switched data transmission system, includes a first group
addressing means, said first group addressing means being adapted to
30 control a multicast to a group address, and to partially resolve a
group address, in that each of a plurality of CLS nodes has a second
group addressing means adapted to complete resolution of a group
address which has been partially resolved by said first group
addressing means, in that only one first group addressing means is
35 permitted per group address and in that more than one second group
addressing means is permitted per group address.

9. A method as claimed in any of claims 6 to 8,
c h a r a c t e r i s e d in that a multicast data packet originating
from a subscriber connected to a first CLS node is transmitted to
said one CLS node, in a data packet which encapsulates said group
5 address.

10. A method as claimed in any of claims 6 to 9,
c h a r a c t e r i s e d in that a multicast data packet transmitted
from a first CLS node and destined for a plurality of subscribers is
10 transmitted from said first CLS node in a data packet which
encapsulates the group address.

11. A method as claimed in any of claims 6 to 10,
c h a r a c t e r i s e d in that said group addressing means, said
15 first group addressing means, and said second group addressing means
are agents, and in that said group addressing means and said first
group addressing means are adapted to interact with a subscriber to
handle requests for the creation, modification, or deletion of group
addresses.

20 12. A method as claimed in any previous claim
c h a r a c e r i s e d in that data packets are transferred along
paths in the form of a tree, packets only being copied where branches
of the tree diverge, and in that the structure of said tree is
25 determined so that system usage is minimised, subject to the
constraint that transmission delays to each group address are kept
within a preset limit.

13. A method as claimed in claim 12, c h a r a c t e r i s e d in
30 that said paths constitute a Steiner tree.

14. A packet switched data transmission system, having a plurality
of switching nodes, each serving a plurality of subscribers,
c h a r a c t e r i s e d in that each node includes copying means
35 for duplicating data packets, intended for transmission to a
plurality of addresses, and transmission means adapted to transmit
copies of multiaddressed data packets to each addressed subscriber
and to the next node to which addressed subscribers are connected.

15. A packet switched data transmission system, as claimed in claim 14, characterised in that said transmission means are adapted to transmit one and only one copy of said data packet to each CLS node;

5

- to which its CLS node is directly connected; and

- to which addressed subscribers are connected.

10 16. A packet switched data transmission system, as claimed in claim 14, characterised in that said transmission means are adapted to transmit one and only one copy of said data packet to each CLS node:

15 - to which its CLS node is directly connected; and

- to which either:

- addressed subscribers; or

20

- a CLS node connected to addressed subscribers;

are connected.

25 17. A packet switched data transmission system, as claimed in any of claims 14 to 16, characterised in that each multicast data packet includes a group address.

30 18. A packet switched data transmission system, as claimed in any of claims 14 to 17, characterised in that it includes at least one ATM based public data network.

35 19. A packet switched data transmission system, as claimed in any of claims 14 to 18, characterised in that at least one CLS node, in said packet switched data transmission system, includes group addressing means, having, for each group address, a data base of all individual addresses comprising a group address, and in that said group addressing means is adapted to control a multicast to a group address and to completely resolve a group address.

20. A packet switched data transmission system, as claimed in claim 18, characterised in that one and only one CLS node, in said packet switched data transmission system, includes group addressing means, having, for each group address, a data base of all individual addresses comprising a group address, said group
5 addressing means adapted to control a multicast to a group address and capable of completely resolving a group address.

21. A packet switched data transmission system, as claimed in claim 10 18, characterised in that one and only one CLS node, in said packet switched data transmission system, includes a first group addressing means, said first group addressing means being adapted to control a multicast to a group address, and to partially resolve a group address, in that each of a plurality of CLS nodes has a second
15 group addressing means adapted to complete resolution of a group address which has been partially resolved by said first group addressing means, in that only one first group addressing means is permitted per group address and in that more than one second group addressing means is permitted per group address..

20

22. A packet switched data transmission system, as claimed in any of claims 19 to 21, characterised in that a multicast data packet originating from a subscriber connected to a first CLS node is transmitted to said one CLS node, in a data packet which
25 encapsulates said group address.

23. A packet switched data transmission system, as claimed in any of claims 19 to 22, characterised in that a multicast data packet transmitted from a first CLS node and destined for a plurality
30 of subscribers is transmitted from said first CLS node in a data packet which encapsulates the group address.

24. A packet switched data transmission system, as claimed in any of claims 19 to 23, characterised in that said group
35 addressing means, said first group addressing means, and said second group addressing means are agents, and in that said group addressing means and said first group addressing means are adapted to interact with a subscriber to handle requests for the creation, modification, or deletion of group addresses.

SUBSTITUTE SHEET

25. A packet switched data transmission system, as claimed in any of claims 14 to 24, characterised in that data packets are transferred along paths in the form of a tree, packets only being copied where branches of the tree diverge, and in that the structure of said tree is determined so that system usage is minimised, subject to the constraint that transmission delays to each group address are kept within a preset limit.

26. A packet switched data transmission system, as claimed in claim 25, characterised in that said paths constitute a Steiner tree.

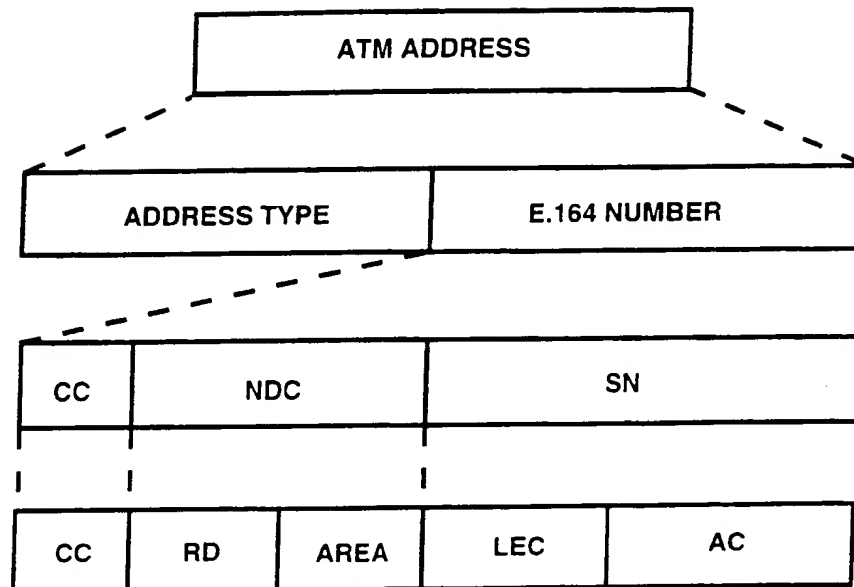


Figure 1




Figure 2

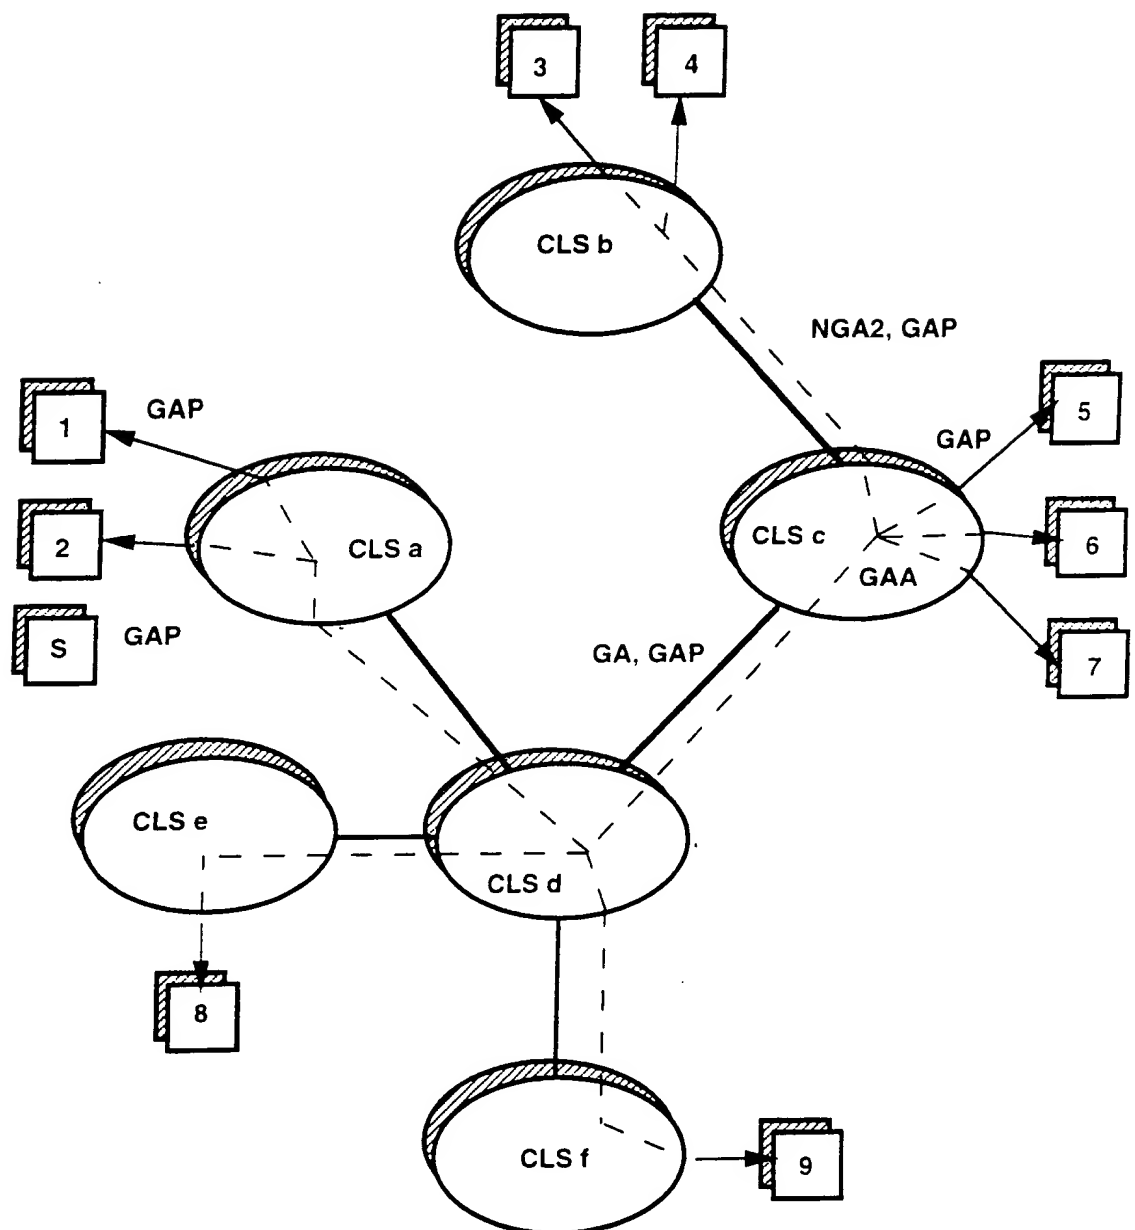


Figure 3

SUBSTITUTE SHEET




Figure 4

5/5

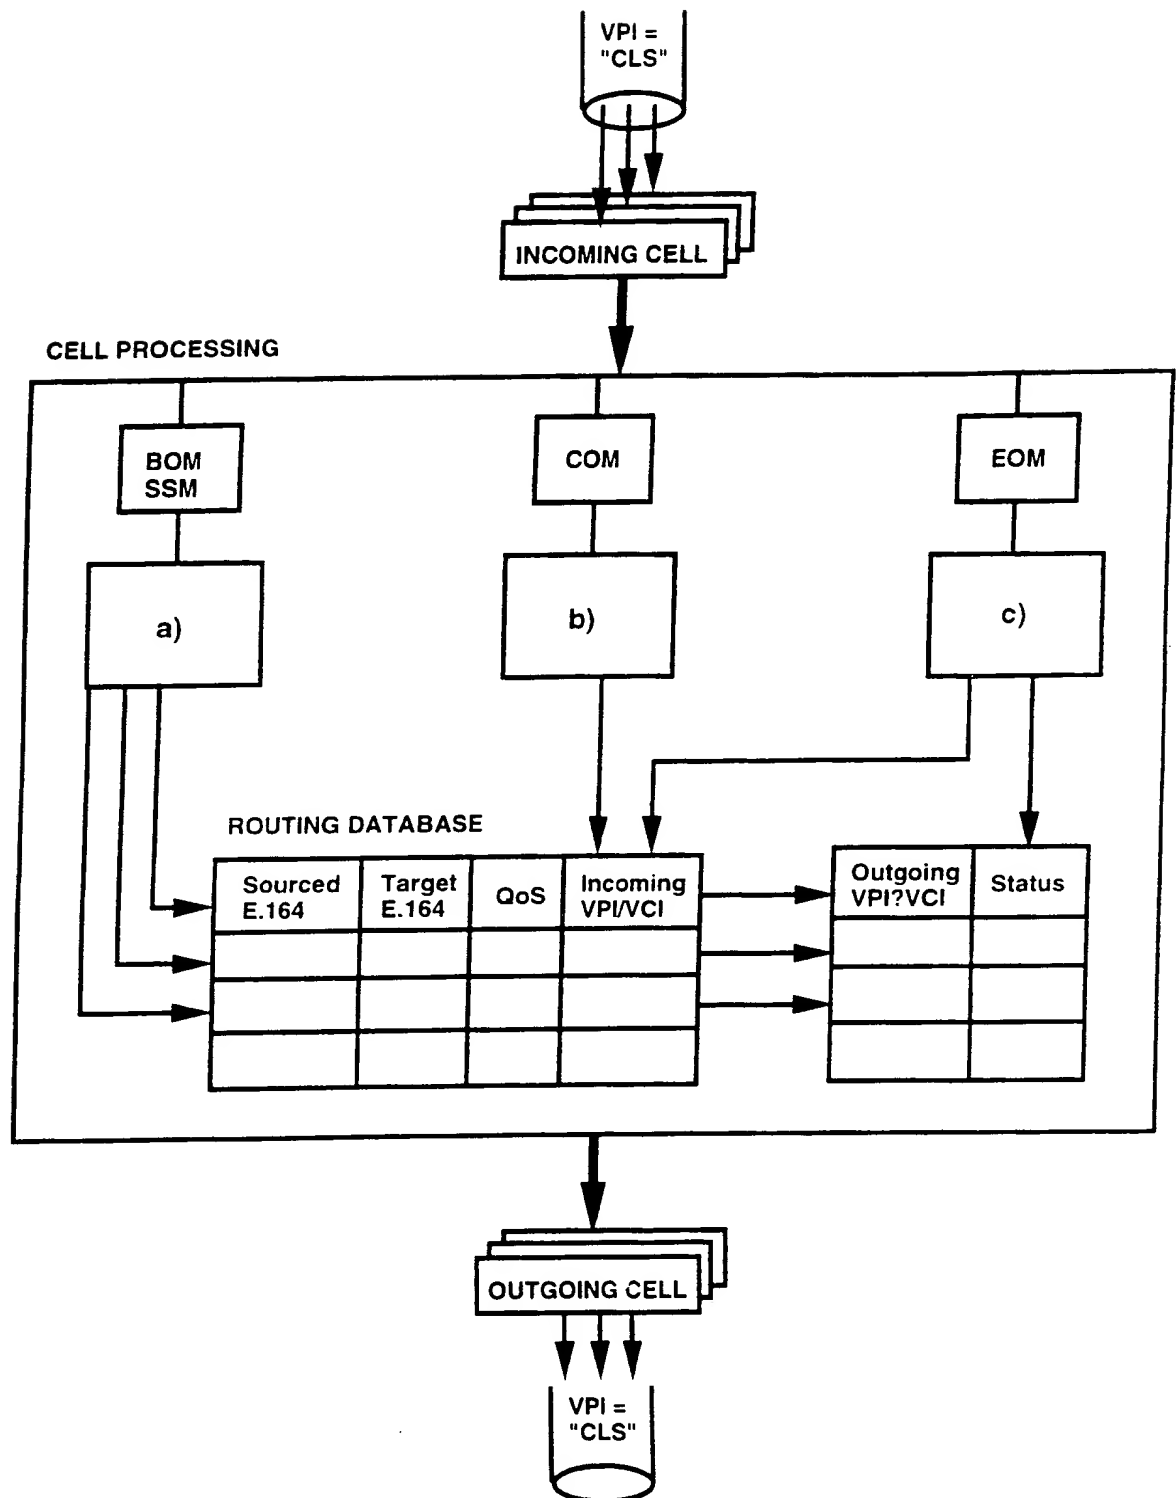


Figure 5

SUBSTITUTE SHEET

INTERNATIONAL SEARCH REPORT

International application No.

PCT/SE 96/00682

A. CLASSIFICATION OF SUBJECT MATTER

IPC6: H04L 12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC6: H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5412649 A (HEINRICH HUMMEL), 2 May 1995 (02.05.95), column 2, line 17 - line 20; column 9, line 23 - line 29; column 10, line 21 - line 35, figures 6,7	1-6,10, 14-19,23
Y		12,13,25,26
A		7-9,11, 20-22,24
	--	
Y	US 5291477 A (SOUNG C. LIEW), 1 March 1994 (01.03.94), column 2, line 65 - column 3, line 14, figure 2C	12,13,25,26
	--	

☒ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"I" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

10 October 1996

16 -10- 1996

Name and mailing address of the ISA/
Swedish Patent Office
Box 5055, S-102 42 STOCKHOLM
Facsimile No. +46 8 666 02 86

Authorized officer

Anders Ströbeck
Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.

PCT/SE 96/00682

C (Continuation). DOCUMENT IS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0637149 A2 (DIGITAL EQUIPMENT CORPORATION), 1 February 1995 (01.02.95), column 3, line 16 - line 29; column 4, line 11 - line 28	1-6,14-19
A	--	11-13,24-26
A	US 4740954 A (CHARLES J. COTTON ET AL), 26 April 1988 (26.04.88), column 2, line 22 - column 3, line 19	1-6,11-19, 24-26
A	--	
A	DE 4330295 A1 (SIEMENS AG), 9 March 1995 (09.03.95), column 2, line 30 - line 66	1-26
	-- -----	

INTERNATIONAL SEARCH REPORT

Information on patent family members

05/09/96

International application No.

PCT/SE 96/00682

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US-A-	5412649	02/05/95	DE-A- 4230744 DE-A,C- 4304120 EP-A- 0589250	11/02/93 15/07/93 30/03/94
US-A-	5291477	01/03/94	NONE	
EP-A2-	0637149	01/02/95	US-A- 5511168	23/04/96
US-A-	4740954	26/04/88	NONE	
DE-A1-	4330295	09/03/95	NONE	

This Page Blank (uspto)